

CLAIMS:

What is claimed is:

1. A cluster multiprocessing system, comprising:
2 a plurality of data processing systems segregated into
3 a plurality of resource groups each including an application
4 and at least two data processing systems;
5 a plurality of configuration objects each corresponding
6 to a resource group within the plurality of resource groups,
7 each configuration object containing:
8 configuration and status information for the
9 corresponding resource group; and
10 an associated owners list identifying data
11 processing systems within the corresponding resource
12 group;

13 a configuration database on each data processing system
14 within the cluster multiprocessing system, each configura-
15 tion database containing at least one configuration object,
16 wherein a configuration object for a resource group is
17 replicated on each data processing system identified in the
18 owners list associated with the configuration object.

1. The cluster multiprocessing system of claim 1, wherein
2 the cluster multiprocessing system comprises N data
3 processing systems and a resource group within the plurality
4 of resource groups contains M data processing systems, where
5 M is smaller than N.

1. The cluster multiprocessing system of claim 2, wherein
2 M equals three and N is at least 1,000.

1. The cluster multiprocessing system of claim 1, wherein

2 each configuration database on a data processing system
3 contains a configuration object for each resource group
4 including the data processing system.

1 5. The cluster multiprocessing system of claim 4, wherein
2 a configuration database on a data processing system
3 included within two resource groups contains two configura-
4 tion objects.

1 6. The cluster multiprocessing system of claim 1, wherein
2 the owners list associated with a configuration object
3 identifies data processing systems which may alter the
4 configuration object.

1 7. The cluster multiprocessing system of claim 1, wherein
2 a configuration object may only be altered by a data
3 processing system identified within the owners list
4 associated with the configuration object.

1 8. A method of managing cluster configuration information,
2 comprising:

3 dividing a cluster into a plurality of resource groups
4 each including an application and at least two data
5 processing systems;

6 instantiating a plurality of configuration objects
7 corresponding to the plurality of resource groups, each
8 configuration object containing:

9 configuration and status information for a
10 corresponding resource group; and

11 an associated owners list identifying data
12 processing systems within the corresponding resource
13 group;

14 maintaining a configuration database on each data
15 processing system within the cluster multiprocessing system,
16 each configuration database containing at least one
17 configuration object,

18 wherein the configuration database on a data processing
19 system contains each configuration object for a resource
20 group which identifies the data processing system as an
21 owner in the owners list associated with the configuration
22 object.

1 9. The method of claim 8, further comprising:

2 replicating a configuration object on each data
3 processing system identified within the owners list
4 associated with the configuration object.

1 10. The method of claim 8, wherein the step of maintaining
2 a configuration database on each data processing system
3 within the cluster multiprocessing system further comprises:

4 maintaining, within the configuration database on a

5 data processing system, a copy of a configuration object for
6 each resource group including the data processing system.

1 11. The method of claim 10, wherein the step of
2 maintaining, within the configuration database on a data
3 processing system, a copy of a configuration object for each
4 resource group including the data processing system further
5 comprises:

6 maintaining a copy of two configuration objects within
7 the configuration database on the data processing system.

1 12. The method of claim 8, wherein the step of
2 instantiating a plurality of configuration objects
3 corresponding to the plurality of resource groups further
4 comprises:

5 listing, within the owners list associated with a
6 configuration object, data processing systems permitted to
7 alter the configuration object.

1 13. The method of claim 8, wherein the step of
2 instantiating a plurality of configuration objects
3 corresponding to the plurality of resource groups further
4 comprises:

5 instantiating one configuration object for a resource
6 group on each data processing system within the resource
7 group.

1 14. A method of partially replicating configuration
2 information in a distributed database, comprising:

3 defining a subset of data processing systems within a
4 cluster system as a resource group;

5 defining configuration data for the resource group; and

6 replicating the configuration data only on each data
7 processing system within the resource group.

1 15. The method of claim 14, wherein the step of defining a
2 subset of data processing systems within a cluster as a
3 resource group further comprises:

4 defining a highly available application and each data
5 processing system designated to manage the application as a
6 resource group.

1 16. The method of claim 15, wherein the step of defining a
2 highly available application and each data processing system
3 managing the application as a resource group further
4 comprises:

5 defining a plurality of resource groups for each highly
6 available application within the cluster, each resource
7 group including all data processing systems managing the
8 corresponding application.

1 17. The method of claim 14, wherein the step of defining
2 configuration data for the resource group further comprises:

3 instantiating a configuration object containing
4 configuration and status information for a highly available
5 application corresponding to the resource group and having
6 an associated list of data processing systems within the
7 resource group.

1 18. The method of claim 17, wherein the step of replicating
2 the configuration data only on each data processing system
3 within the resource group further comprises:

4 replicating the configuration object on each data
5 processing system identified in an owners list associated
6 with the configuration object.

1 19. The method of claim 17, wherein the step of replicating
2 the configuration data only on each data processing system
3 within the resource group further comprises:

4 replicating, on a data processing system, a configura-
5 tion object for each resource group including the data
6 processing system.

1 20. The method of claim 14, further comprising:

2 maintaining, on a data processing system, a configura-
3 tion object for each resource group including the data
4 processing system and no configuration objects for other
5 resource groups.

1 21. A data processing system, comprising:
2 a processor executing instructions for an application
3 server;
4 a connection port permitting connection of the data
5 processing system to a cluster system to receive request for
6 the application server; and
7 a memory containing configuration information for the
8 cluster system and configuration information for at least
9 one resource group within the cluster system, wherein the
10 resource group is associated with the application server,
11 the configuration information for the at least one resource
12 group including an identification of other data processing
13 systems in the resource group.

1 22. The data processing system of claim 21, further
2 comprising:

3 means, responsive to a configuration change, for
4 determining whether the configuration change affects the
5 cluster system or only the resource group;

6 means, responsive to determining that the configuration
7 change affects the cluster system, for replicating change
8 information to all data processing systems in the cluster
9 system; and

10 means, responsive to determining that the configuration
11 change affects only the resource group, for replicating
12 change information only to data processing systems within
13 the resource group.

1 23. A computer program product in a computer usable medium,
2 comprising:

3 instructions defining a subset of data processing
4 systems within a network as a resource group;

5 instructions defining configuration data for the
6 resource group; and

7 instructions for replicating the configuration data
8 only on each data processing system within the resource
9 group.

24 23. The computer program product of claim 22, wherein the
instructions defining a subset of data processing systems
within a network as a resource group further comprise:

 instructions defining a highly available application
and each data processing system designated to manage the
application as a resource group.

25 24. The computer program product of claim 23, wherein the
instructions defining a highly available application and
each data processing system managing the application as a
resource group further comprise:

 instructions defining a plurality of resource groups
for each highly available application within the network,
each resource group including all data processing systems
managing the corresponding application.

26 25. The computer program product of claim 22, wherein the
instructions defining configuration data for the resource
group further comprises:

 instructions instantiating a configuration object
containing configuration and status information for a highly
available application corresponding to the resource group

7 and having an associated list of data processing systems
8 within the resource group.